

1        **COMPONENTIZED AUTOMATIC PROVISIONING AND MANAGEMENT OF**  
2        **COMPUTING ENVIRONMENTS FOR COMPUTING UTILITIES**

3        **CROSS REFERENCES**

4        The present application is cross-referenced to application entitled "Hierarchical Resource  
5        Management for a Computing Utility", having docket number YOR920030587PCT1, and  
6        application entitled "Apparatus for Arbitration in a Computing Utility System", having docket  
7        number YOR920030589US1, even dated, and which are included herein by reference in entirety  
8        for all purposes.

9        **FIELD OF THE INVENTION**

10       The present invention is directed to provisioning and management of resources in a computing  
11       utility, which uses the resources to provide computing services to customers. It is more  
12       particularly directed to componentized automatic provisioning and management of computing  
13       environments for the computing utility.

14       **BACKGROUND**

15       A hosting center provides computing services to multiple customers. Each customer is allocated  
16       a subset of the service provider's infrastructure resources, such as servers, to meet its needs.  
17       Customer needs change over time, and in particular, peak demands for resources can exceed  
18       average demands by orders of magnitude. A simple static allocation of resources to satisfy peak  
19       demand leaves resources underutilized. Reconfiguring infrastructure resources dynamically, in

response to customer needs, demands prompt attention from administrative personnel, increasing both operational costs and the risk of failing to provide adequate service. The problem for the hosting center is how to respond quickly to changes in customer needs such that infrastructure resources and staff are employed in an efficient and cost effective way. *Computing utilities* attempt to address this problem by automating the creation and management of one or more computing services on shared infrastructures.

Previous work in computing utilities varies in the types of services that are offered, the resources that are used, and the extent to which operation is automated. Operations subject to automation are wide ranging, and include creating services, deploying a service for a customer, modifying the set of resources used to provide the service, and incorporating new resource instances and types into the hosting center and its services.

One system automated the provisioning of front end servers in Web sites based on metrics such as server load and response time. It included a component that discovered server and network topology automatically. Another system provided a variety of multi-tier Web sites in which pre-configured servers could be allocated to different tiers automatically based on metrics such as server load. Another system also allocated server resources in response to server load, but modeled both the value of allocating resources to each customer and the cost of employing those resources, with an emphasis on energy cost. More recent work includes allocation of other resource types such as memory and storage, and allocation of servers for general use.

A growing number of industrial products aim to provide multi-tier applications over a physical infrastructure consisting of a variety of resource. They vary in many respects, such as the types of resources provided (e.g., servers and storage); specific operating systems and middleware supported; assumptions and characteristics of the network infrastructure (e.g., whether or not network isolation is provided via VLAN); level of monitoring support (e.g., resource usage, failure detection, SLA, threshold based alerting); support for resource discovery; support for modifying service resources once deployed; whether modifications can occur automatically (e.g.,

1 triggered by SLAs); and the extent to which the products can or must be customized to fit  
2 preexisting hosting center infrastructures.

3 Work is being done at the application layer to automate deployment and management of  
4 distributed applications. An application is described as a collection of related, reusable  
5 components, which may represent resources or subsystems. The description of an application  
6 includes dependency information to ensure that, for example, components are started in the  
7 correct sequence. This is not used for low level resource configuration tasks such as installing  
8 operating systems on servers, but for higher level application specific configuration.

9 Provisioning refers to the process of allocating a set of infrastructure resources and configuring  
10 them to provide the desired computing services for use by a customer. Provisioning is a  
11 complex set of tasks that can include configuring network devices, installing operating systems  
12 and application software, obtaining software licenses, installing monitors, and enabling customer  
13 access by creating user accounts.

14 Provisioning is automated by codifying the procedures or scripts for these tasks. Some products  
15 provide a GUI interface to assist in the creation and selection of these procedures. For example, a  
16 drag and drop interface may be used to specify the “virtual wiring” of resources for an  
17 application environment to be deployed. Virtual wiring could include the tier structure of the  
18 environment, rules for firewalls, and which operating system to install on servers. The wiring  
19 diagram is captured as a template description. This description is used to create and delete  
20 environments, and it may be modified to add or remove resources from a computing  
21 environment.

22 In some systems codifying the procedures is done on a service specific basis. Changes in resource  
23 types, connectivity, or service requirements can require re-implementation of these procedures.  
24 If they are implemented in an ad hoc manner or independently, it can be difficult to

1 systematically enforce best practices or operational constraints. Thus, it is desirable to provide a  
2 system that improves upon the prior art.

### 3 **SUMMARY OF INVENTION**

4 This invention provides methods and apparatus for automatically provisioning and managing  
5 resources in a computing utility. Its automation procedures are based on a resource model which  
6 allows resource-specific provisioning and management tasks to be encapsulated into components  
7 for reuse. These components are assembled into more complex structures and finally computing  
8 services. Included are methods and apparatus for automatically provisioning and managing  
9 resources in a computing utility. Its automation procedures are based on a resource model which  
10 allows resource-specific provisioning and management tasks to be encapsulated into components  
11 for reuse. These components are assembled into more complex structures and finally computing  
12 services. This invention also provides a method for constructing a computing service from a set  
13 of resources given a high level specification. Once constructed, the service includes a component  
14 that provides management function, which can allow modification of its underlying set of  
15 resources.

16 In an example embodiment, this invention provides a method for provisioning at least one  
17 computing environment in a computing utility. The method comprising: employing a description  
18 of each of said at least one computing environment, said description directing an invocation of a  
19 first service which performs said provisioning, said description being a composite resource  
20 definition; and invoking said first service to perform said provisioning.

21 In an example embodiment, this invention provides a method for adding a base resource type to  
22 an operating computing utility. The step of adding comprising: furnishing a base resource library  
23 service to represent any instance of said base resource type; adding any instance of said base

1 resource type to the base resource library service; developing an implementation of a  
2 provisioning resource service for said base resource type; activating said provisioning resource  
3 service; and creating an implementation of a managed resource service for said base resource  
4 type.

5 In some embodiments, the method further comprising: updating any provisioning resource  
6 service of any composite resource type to use said provisioning resource service for said base  
7 resource type; and updating any managed resource service of any composite resource type to use  
8 said managed resource service for said base resource type.

9 In a further example embodiment, this invention provides a method for adapting at least one third  
10 party workload management system for a computing environment to a computing utility. The  
11 step of adapting comprising: developing a managed resource service, said managed resource  
12 service interpreting a state of said at least one third party workload management system; and  
13 modifying a set of base resources comprising said computing environment as indicated by the  
14 state of said at least one third party workload management system.

15 In an further example embodiment, this invention provides a method for constructing a  
16 computing service from a set of resources given a high level specification. Thus, this invention is  
17 also an element of a computing utility.

18 In a still further example embodiment, this invention provides an apparatus to provision at least  
19 one computing environment in a computing utility. The apparatus comprising: a first processor  
20 module to employ a description of each of said at least one computing environment, said  
21 description directing an invocation of a first service which performs said provisioning, said  
22 description being a composite resource definition; and a second processor module to invoke said  
23 first service to perform said provisioning.

## **BRIEF DESCRIPTION OF THE DRAWINGS**

The foregoing and other objects, aspects and advantages will be better understood from the following detailed description of advantageous embodiments of the invention with reference to the drawings, in which:

Fig. 1 depicts a hosting center, a system in which the present invention may be used;

Fig. 2 shows an example of a computing utility structure, in which the PMRS, in accordance with the present invention, is one of the components;

Fig. 2a shows an example of an internal structure of the PMRS in accordance with the present invention;

Fig. 3 shows a composite resource instance for a Web site;

Fig. 4 shows an MRS hierarchy that corresponds to the composite resource instance of Fig. 2 in accordance with the present invention;

Fig. 5 shows a PRS hierarchy that corresponds to the composite resource instance of Fig. 2 in accordance with the present invention;

Fig. 6 shows a realization for a Web site composite resource in accordance with the present invention;

Fig. 7 shows a Web site realization pruned by tier attribute in accordance with the present invention;

Fig. 8 shows a Web site realization pruned by server attribute in accordance with the present invention;

Fig. 9 shows a Web site realization pruned by tier, Server, and, WAS attributes in accordance with the present invention;

Fig. 10 shows a Web site realization and its components, where each component corresponds to a PRS and MRS structure in accordance with the present invention;

Fig. 11 shows an example of operations of the Realizers involved in the construction of a Web site realization in accordance with the present invention;

Fig. 12 shows a process used to construct a composite resource in accordance with the

present invention;

Fig. 13 shows a process used to construct a composite resource in a system with advanced reservations in accordance with the present invention;

Fig. 14 shows an example of PMRS interaction with BRDS in accordance with the present invention;

Fig. 15 shows an example of MRS interaction with a third party workload management system in accordance with the present invention;

Fig. 16 shows s an example of how to add a new base or composite resource type to the system in accordance with the present invention;

Fig. 17 is a use case that describes adding a Linux server to a Linux server group service based on a load event in accordance with the present invention; and

Fig. 18 is a use case that describes adding a Web server to a Web site service based on a load event in accordance with the present invention.

## **DEFINITION OF TERMS**

A **computing environment** is a computing service provided to a customer. It includes hardware and software resources such as servers, operating systems, and middleware. This invention assumes that the resources allocated to a computing environment are dedicated to that computing environment for the duration of their allocation.

A **service provider** is assumed to be providing service for multiple computing environments. A service provider supplies a set of **resources** that may be allocated to provide computing services to its computing environments.

A **base** resource is a resource that is atomic; it cannot be broken down into other resources. An example of a base resource is a server or a software license. Any resource that a service provider does not want to subdivide can be treated as a base resource. Adding, removing, or modifying a

1 base resource changes the overall capacity of the system. We assume that both the number and  
2 types of the base resources change over time. The notion of base resource is service provider  
3 specific.

4 A **composite** resource is built out of other resources to perform a specified function. A  
5 composite resource may have a set of **dependencies** associated with it. An example of a  
6 composite resource is a Web site. It may be composed of some number of front-end servers, back  
7 end servers, a load balancer, a set of IP addresses for the servers, Web server software, database  
8 software, and licenses associated with the software. The set of base resources used to implement  
9 the function of a composite resource can change over time, though not all composite resources  
10 may have this capability. In some cases, a resource within a composite resource is referred to as  
11 a **sub-resource**.

12 A resource may be requested by a computing environment either for an immediate or future  
13 allocation. The promise of a future resource allocation is a **reservation**.

## 14 DESCRIPTION OF THE INVENTION

15 This invention provides methods, systems and apparatus for automatically provisioning and  
16 managing resources in a computing utility in an environment that possesses the characteristics  
17 listed below. Together, these characteristics generalize the environments studied in previous  
18 work.

19 The present invention provides methods and apparatus for automatically provisioning and  
20 managing resources in a computing utility. Its automation procedures are based on a resource  
21 model which allows resource-specific provisioning and management tasks to be encapsulated  
22 into components for reuse. These components are assembled into more complex structures and  
23 finally computing services. This invention provides a method for constructing a computing  
24 service from a set of resources given a high level specification. Once constructed, the service



includes a component that provides management function, which can allow modification of its underlying set of resources.

An example method provides for constructing a computing environment from a set of resources given a high level specification. It is a component of a computing utility that operates in an environment that possesses at least one of the characteristics listed below. Together, these characteristics generalize the environments studied in previous work.

First, resources may be allocated to customers in combinations which are heterogeneous, may be interdependent, and vary over time.

Second, the types of computing environments provided to each customer may be different. For example, one customer may be provided resources for a Web site, and another for a scientific computing cluster. Resource types, quantities, dependencies, and allocation patterns will thus vary between customers.

Third, there can be multiple ways to construct a computing environment from the resources in a service provider's infrastructure. A customer may have preferences or requirements for particular variations of a given computing environment, such as for particular brands of servers.

Fourth, the resource infrastructure varies between service providers. Further, for a given service provider, the infrastructure varies over time. These variations can be a result of upgrades or additions to the physical infrastructure, engaging the services of another provider, or inclusion of additional resources when idle or at certain times of day.

Finally, resources may be allocated in advance, or reserved.

1 The system in which the present invention is used is a distributed computer system which has  
2 multiple computing resources interconnected via a network. A diagram of such a system, which  
3 would be found in a hosting or data center, is shown in Figure 1. The computing resources in the  
4 physical infrastructure include processors 101, storage 103, firewalls 105, and software 107.

5 The software can be operating systems, middleware or applications. In figure 1 the available  
6 software has been pre-configured into business processes, business applications, or services. The  
7 hardware resources are connected by a network 109 as indicated by the grid of lines  
8 interconnecting all of these resources. This network may be configure into one or more tiers,  
9 where each tier is separated by a router or firewall. Software resources are assigned to physical  
10 resources by a management infrastructure. In this environment a subset of the resources are  
11 assigned to the management infrastructure. In Figure 1 these resources 111, 113, and 115 are  
12 indicated by a dotted line surrounding them. The resources assigned to the management  
13 infrastructure run the management software described in this invention. This software manages  
14 the rest of the resources. The resources used for the management infrastructure are not assigned  
15 to customers of the hosted environment. The remaining resources are assigned to customers as  
16 necessary. It is expected that clients 117 will primarily receive service by connecting to their  
17 resources through the Internet 119. However, clients can receive service if they are connected to  
18 their resources by any means, such as through a direct connection to a managed resource.

19 Figure 2 shows a high level view of a computing utility in which the present invention is a  
20 component. There are three layers of components. The bottom layer 201 represents the base  
21 resources. Each base resource type has a service that provides reservation and allocation of  
22 instances of that resource type, called the Base Resource Library Service (BRLS). This invention  
23 assumes that both the number and types of the base resources change over time. Figure 2 shows  
24 four types of base resources and their BRLSs: DB2 software licenses 203, zSeries logical  
25 partitions 205, xSeries server hardware 207, and AIX licenses 209. The BRLS provides  
26 information such as cataloging, check out (allocation), check in (deallocation), and reservation.  
27 Each computing environment has a service which is invoked to provision and manage its  
28 resources, called the Provisioned and Managed Resource Service (PMRS). In Figure 2, the top

layer 211 shows the computing environments and their PMRSs. Figure 2 shows three examples of computing environments, a DB2 environment 213, a Linux environment 215, and a Web site environment 217. In the middle layer, the Base Resource Distribution Service (BRDS) 219 determines how to distribute the base resources across the various computing environments. The BRDS specifies which sets of resources are available to which computing environments and automatically redistributes those resources based on business policies when resources are constrained.

The resources assembled for a computing environment are modeled as composite resources. The model defines the structure of the composite resource in terms of its constituent resources. A set of services, called Provisioned and Managed Resource Services (PMRS) which are the focus of this invention, provision and manage composite resources. Every resource type has a PMRS that encapsulates the knowledge of how to create a resource of that type, and how to manage an instance of a resource of that type. Managing the resource may include such actions as adding or removing capacity from the resource instance. Thus, as shown in Figure 2a, a PMRS 251 is composed of two parts: a Provisioning Resource Service (PRS) 253, which creates instances of the composite resource, and a Managed Resource Service (MRS) 255, which manages an instance of the composite resource. There is a single PRS for each resource type, and there is an MRS for every instance of a resource type.

Composite resources may be assembled using other composite resources. For example, a Web site composite resource instance may be composed of a load balancer, and a group of front end servers. A front end server is a composite resource which may be composed of a server, operating system, and Web server software. For this reason, PMRSs can be composed in a containment hierarchy that represents the structure of the composite resources in the computing environment. Figure 3 shows an example of a Web site composite resource instance. Every node in the tree represents either a basic or a composite resource. The root of the tree 301 is the Web site composite resource. It contains a load balancer resource 303, which is a basic resource, and a Web server group composite resource 305. The Web server resource in turn contains a server

resource 307.

A PMRS of a composite resource may use the PMRS of a sub-resource to create and manage instances of the sub-resource. This approach allows resource provisioning and management to be encapsulated on a per resource type basis. A PMRS for a composite resource only has to incorporate and reflect knowledge of the interdependencies between the sub-resources of the composite, since the internal knowledge of a sub-resource is encapsulated in the sub-resource's PMRS.

Figure 4 shows the MRS hierarchy which corresponds to the Web site composite resource instance 301. Every MRS is responsible to manage a single resource instance. Thus, a Web-site MRS 401 is responsible to manage the entire Web site, whilst a Web server MRS 403 is responsible for managing Web server 309. An arrow from an MRS box to another indicates that the former uses the latter in order to perform its function. Thus, a Web-site MRS 401 may use information provided by the load balancer MRS 405 and the various Web server MRSs 407 in order to determine the root cause of a problem in the performance of the Web site, and web server MRS may in turn use server MRS 409. Web site MRS 401 may further use the MRSs of the sub-resources in order to perform corrective configuration actions on some composed resources. The process will be further described later.

Figure 5 shows the PRS hierarchy which corresponds to the Web site composite resource instance from Figure 3. A PRS in this figure is responsible to create all instances of resource from the corresponding type. Thus, the root of the hierarchy is a Web site PRS 501 which is responsible for creating all Web site composite resource instances in the computing utility. The Web server PRS 503 is responsible for creating all Web server instances in the computing utility and so on. As in the case of the MRS hierarchy, an arrow from a PRS box to another indicate the former uses the latter in order to perform its function. Thus, the Web site PRS 501 uses both the Web server PRS 503 and the load balancer PRS 505 in order to build a Web site resource. It does not need to know how to build a Web server or a load balancer; this knowledge is encapsulated

1 in the corresponding PRSs that it uses. It only needs to know that a Web site consists of Web  
2 servers and a load balancer and how to configure them, once they are created, to work together.  
3 The process will be further described later.

4 Figure 3, Figure 4, and Figure 5 illustrate the idea of applying a recursive structure to the  
5 provisioning and management of systems. Thus, PRSs, and MRSs serve as generic building  
6 blocks for building more complex resources. An implementer of a PRS or MRS for a new  
7 composite resource does not have to re-implement the logic to create and manage any of the  
8 sub-resources. The implementer only needs to implement the increment; how the resources can  
9 be created and managed to work together to perform a higher function beyond the function that  
10 any single resource performs.

## 11 **Provisioned Resource Service**

12 A PRS creates instances of a composite resource type. It implements two phases in the creation  
13 of a composite resource – a planning phase and a building phase. In the PRS, the planning phase  
14 is implemented by a component called the *realizer*. The realizer maps an abstract description of  
15 a resource to a schema that describes the possible combinations of base resources needed to  
16 create it and the corresponding composite resource structure. It essentially supplies a plan for  
17 building the composite resource. Since there may be more than one way to create the resource,  
18 the realizer can supply multiple plans. The realizer takes into account resource types, but not  
19 availability of instances of a particular type, when providing plans. A realizer may use the  
20 realizers in PRSs of sub-resources, to create plans.

21 A second component, the *builder*, builds a resource instance given resource handles for the base  
22 resources required and a plan. It also disassembles resources when requested. Creating an  
23 instance of a composite resource may be multi-staged, where at each stage, resources are built  
24 from sub-resources incrementally by calling the Builder services of the sub-resources  
25 provisioning services, and then configured with respect to each other. The creation of a resource

instance includes the instantiation of an MRS to manage it. Following the same logic, in the process of a creation of a composite resource a set of MRSs will be instantiated to manage the composite resource and its sub-resources cooperatively.

#### **Realizer**

The realizer encapsulates the possibilities for constructing a resource of a particular type from the base and composite resources available in the computing utility. It may supply one or more alternative plans for constructing the resource. In this case, it may also expose certain settings that can be used to restrict the set of alternatives it supplies. A realization is a structure that encodes the different alternatives. The realization can be represented in a number of structures such as a graph or binary decision diagram. The representation should be compact because the possible number of realizations of a complex structure could be large.

An example realization for a Web site resource is shown in Figure 6. The nodes in the graph with children are of two types: AND nodes, shown in black, and OR nodes, shown in white. An AND node means that all the children of the node are required to construct the resource represented by the parent. An OR node means that any of its children may be used to construct the resource represented by the parent. The root node of the Web site graph 601 is an OR node, representing a two-tier alternative 603 and a three-tier alternative 605. The two-tier node 603 is an AND node composed of a load balancer 607 and a Web server 609. The three-tier node 605 is also an AND node composed of a load balancer 611, front end Web server 613, and back end Web server 615. The load balancer nodes 607 and 611 are base resources; they have no outgoing edges. The two-tier Web server node 609 is an OR node, with only one alternative; server node 617. The server node 617 is an OR node representing alternatives Apache Web server 619 and WAS Web server 621. The three-tier front end Web server node 613 is identical in structure to the two-tier Web server node 609. Apache Web server node 619 is an OR node representing servers of three IBM architectures, iSeries 623, xSeries 625, and pSeries 627. WAS Web server 621 is an OR node representing servers of four IBM architectures, iSeries 629, xSeries 631,

1 pSeries 633, and zSeries 635. These servers are base resources, as is DB2 node 637.

2 A realizer for a resource type may allow the setting of certain attributes that can restrict the set of  
3 alternatives it supplies. These settings are specified as part of the realization-spec parameter of  
4 the realize() function described above. In the AND-OR graph representation shown in Figure 7,  
5 an attribute 701 specifies a condition on the value of the descendant that is chosen for an OR  
6 node. The graph is pruned to eliminate descendants that do not satisfy the condition in the  
7 attribute. Figure 7 shows the possible realizations corresponding to the “web site is two-tier”  
8 attribute. This attribute, if set, prunes the graph at the “web site” OR node 601. The resulting  
9 realization is the sub-tree rooted at node 603, shown with solid edges.

10 A more complex example is shown in Figure 8. The attribute “Server is Apache” 801 specifies  
11 that any “server” OR node in the realization must have an “Apache” node as a direct descendant,  
12 such as the descendant of server node 617. Other descendants of the “server” OR node are  
13 pruned, such as node 621. If an “Apache” node is not specified as a direct descendant of a  
14 “server” node in the initial realization, as in node 807, the entire “server” node is pruned. If the  
15 pruned “server” node is the only direct descendant of an OR node, or if it is any direct descendant  
16 of an AND node then this node is pruned, such as node 615, and so on in a recursive manner,  
17 with node 605. The application of this process results in the realization that is the two-tier  
18 realization rooted at node 603, with a web server 617 realized as Apache server 619. The  
19 realization is shown as with solid arrows. Any descendant of node 619 is an acceptable  
20 alternative.

21 Multiple attributes may be specified as part of the realization-spec. Figure 9 shows a  
22 realization-spec with multiple attributes 901. The resulting realization is three-tier web site 605  
23 with front end server 905 realized as WAS server 907 on server type zSeries 909. Similarly, back  
24 end web server realized as WAS server 911 is also realized with zSeries 913. This realization is  
25 shown with solid arrows.

1 The above pruning algorithm can be semi-formally described as follows:

2 **Prune\_one (S,R):**

3 For an attribute  $S = [ \text{NODE\_NAME is } \textit{CONDITION} ],$

4 and for a realization  $R$  with a set of nodes  $\textit{NODES}.$

5     1. Let  $\textit{Pruned} = \textit{empty-set}.$

6     2. For every node  $N \in \textit{NODES}:$

7         If  $N$  is an **OR** node then

8             If  $N = \text{NODE\_NAME}$  then

9                 For every direct descendant  $M$  of  $N$  in  $R:$

10                     If  $M$  does not satisfy  $\textit{CONDITION}$  then

11                          $\textit{Pruned} = \textit{pruned} \cup \{M\}$

12     3. For every node  $N \in \textit{NODES} \setminus \textit{Pruned} :$

13         If  $N$  is an **OR** node then

14             If all descendants of  $N$  are in  $\textit{Pruned}$  then

15                  $\textit{Pruned} = \textit{pruned} \cup \{N\}$

16     4. For every node  $N \in \textit{NODES} \setminus \textit{Pruned}:$

17         If  $N$  is an **AND** node then

18             If  $N$  has a descendant  $M$  s.t.  $M \in \textit{Pruned}$  then

19                  $\textit{Pruned} = \textit{pruned} \cup \{N\}$

20     5. Repeat steps 3 and 4 until an iteration is reached where the set  $\textit{Pruned}$  is not changed.

21     6. Return the tree of nodes that are reachable from the root and are not pruned (if the root

22     is pruned return the empty tree).

23 **Prune\_all(LS,R):**

24 For a list of attributes  $LS = L1, L2, \dots, Ln,$

25 And for a realization  $R$



```

1      1. New_R = R
2      2. For i=1,...,n:
3          New_R = prune_one(New_R,Li)
4      3. Return New_R

```

Note that these examples are based on a realizer that allows all of the OR node attributes to be specified in the realize request. The realizer need not expose all of these choices. It may also group a set of choices as a single attribute, such as “non-Intel servers”.

Also note that more complex attributes may be defined and used. For example, using the keyword *contains* instead of *is*, attributes can be specified on nodes in the sub-tree rooted by a node and not only on direct descendant of it. The pruning algorithm can be changed in a straight forward way to handle also these cases. As an example, consider again the result of the realization described in Figure 9. The same realization result also corresponds to the following realization-spec with only two attributes: “Web-site *is* two-tier” and “Server *contains* z”. Although it is not specified that a server must be WAS, the Apache possibility 915 is pruned because the sub-tree rooted by it does not contain a “z” resource.

## Operation of the Realizers

Since composite resources may be assembled using base or other composite resources, realizers invoke the realizers of sub-resources. Figure 10 shows the Web site example using two sub-resources, Web server and server. The Web server realizer supplies the sub-graphs rooted at striped nodes 1001, 1003, and 1005. This realizer can supply front end or back end web server types. Node 1005 is darker than nodes 1001 and 1003 because it is both an AND node and a sub-resource node. The server realizer supplies the sub-graphs rooted at dotted nodes, such as node 1007. The server realizer can supply servers of the different IBM server architectures with compatible operating systems. Note that in this example, the Apache Web server realizer calls the server realizer 1007 with only the iSeries, xSeries, or pSeries server type. The zSeries option

is not shown for Apache in the graph.

Figure 11 shows the process of generating the realization. The Web site PRS 501 receives Realize request 1103 with a realization-spec with the attribute “two – tier”. Web site realizer 501 calls the web server realizer 503 with Realize request 1107. The web server realizer 503 in turn calls the server realizer 507 twice, with Realize requests 1111 and 1113. The results of requests 1111 and 1113 are the realizations **R1** 1115 and **R2** 1117 returned by responses 1119 and 1121. The Web server realizer uses **R1** and **R2** to build the realization **R3** 1123 which is returned on response 1125 in answer to request 1107. Last, the Web-site realizer user **R3** to construct the realization **R4** 1127 which is returned on response 1129 from the initial request 1103.

Realizations may be obtained well in advance of the time that a composite resource is constructed, and they may be retained for long periods. Circumstances under which realizations should be regenerated include changes in requirements and changes in the base or composite resource types available in the system. Note that a realization does not depend on resource availability. Therefore, changes in the availability of resources do not require regeneration of realizations.

## **Resource Selection**

Once a realization is obtained, the resources corresponding to a single alternative must be allocated. Since a realization can describe multiple alternatives, a ranking mechanism can be used to order the alternatives according to some metric. For example, a set of preferences can be indicated for the remaining alternatives within the realization, and the alternatives could be ordered in terms of number of preferences satisfied. Another ranking metric could be cost. The alternatives are then selected in turn and the base resources are requested for a specific alternative. If the resource request is successful, the tree is annotated with the resource identifiers of the base resources (e.g., reservation tickets or resource handles). Once base resources are obtained, the builder can be invoked to construct the resource. Note that if resource

availability changes (e.g., a reservation is canceled by the system), the realization can be used to secure an alternative set of resources.

### **Builder**

The builder receives as input a realization and a set of base resources from which to assemble the resource. The procedures needed for construction of the resource type are encoded in the builder. These procedures encapsulate details of resource dependencies and proper ordering of operations. The realization is annotated with the alternative corresponding to the base resources. The builder traverses the tree in depth-first manner, invokes the builders of sub-resources at nodes as indicated and supplies them with the sub tree of the realization that corresponds to the sub resource.

Figure 12 shows the process of constructing a composite resource. The process begins in step 1201, in which the realize function of the PRS for the composite resource is called with requirements supplied as attribute settings in the realization-spec, if any. Next in step 1203 the realizations are ranked and ordered by rank. If there are preferences over the alternatives they are addressed in this step. The process then examines the alternatives embodied by the realizations in turn, starting in step 1205 with the first realization. The set of base resources is obtained for the realization in step 1207, and a resource request is issued for the set in step 1209. The result of the request is examined in step 1211. If the request succeeds, flow passes to step 1217, in which the realization is annotated with the current alternative and the resource-ids of the set of base resources secured for it. Then in step 1219 the build function of the PRS is invoked. If in step 1211 the request was not successful, flow passes to step 1213 to determine if there is an alternative realization that may be tried. If there is an alternative realization, the next realization on the list is selected in step 1215 and the process continues with step 1207. If there are no alternatives left, flow passes to step 1221 and the construction of the composite resource fails.

In a system that supports advance reservations, the resource identifiers returned may represent

1 reservations for the resource. In that case, another step is needed to claim the resources  
2 associated with the reservation. There may be a considerable period before this step occurs, and  
3 as a result of other activities during that period, this step may eventually fail. For example, the  
4 resources may not be available at the time the reservation is due to begin, for reasons such as  
5 failure of the resource, or overbooking of reservations. If the initial attempt to claim the resources  
6 is not successful, other alternatives can be tried and recovery may be possible with another set of  
7 base resources. The process of constructing a composite resource in a system that supports  
8 advance reservations is shown in Figure 13. The process begins in step 1301, in which the  
9 realize function of the PRS for the composite resource is called with requirements supplied as  
10 attribute settings in the realization-spec, if any. Next in step 1303 the realizations are ranked and  
11 ordered by rank. If there are preferences over the alternatives they are addressed in this step.  
12 The process then examines the alternatives embodied by the realizations in turn, starting in step  
13 1305 with the first realization. The set of base resources is obtained for the realization in step  
14 1307, and a resource request is issued for the set in step 1309. The result of the request is  
15 examined in step 1311. If the request succeeds, flow passes to step 1317, in which the realization  
16 is annotated with the current alternative and the resource identifiers of the set of base resources  
17 reserved for it. When the construction time arrives in step 1318, the process attempts to claim  
18 the resources associated with the reservation in step 1319. The construction time allows for  
19 sufficient time to construct the resource so that it is available at the desired start time. If all of the  
20 resources are acquired in step 1321, flow passes to step 1323 where the build function of the PRS  
21 is invoked. If all of the resources are not acquired in step 1321, control passes back to step 1305  
22 where the process is repeated starting with the first realization. If in step 1311 the resource  
23 request was not successful, flow passes to step 1313 to determine if there is an alternative  
24 realization that may be tried. If there is an alternative realization, the next realization on the list  
25 is selected in step 1315 and the process continues with step 1307. If there are no alternatives left,  
26 flow passes to step 1325 and the construction of the composite resource fails.

## 27 *Provisioning Operations*

1       **Realize(n, realization-spec) -> realization[]**

2           This request is issued to create a plan for building a resource of the type  
3           represented by the PMRS.

4           Inputs:

5               n, the maximum number of realizations desired

6               realization-spec, the settings for attributes exposed by this realizer

7                       These attributes control represent possible variations in the  
8                       resource. Their settings determine the realizations supplied

9                       Unspecified attribute settings represent additional degrees of  
10                      freedom.

11          Outputs:

12               Realization[]. A structure representing the realizations. (list, tree, etc.).

13               Each realization includes a schema that contains the base resources  
14               required for it (“shopping list”) and their structure.

15       **Build(base-resource-handles, r) -> resource-handle**

16           This interface is used to request the build of a resource of the type represented by  
17           the PMRS from the base resources indicated by base-resource-handles.

18           Inputs:

19               Base-resource-handles[], a list handles for the base resources in the  
20               realization

21               Realization, provided previously by the realizer that is used to describe the  
22               structure of the composite resource.

23          Outputs:

24               Resource-handle, a handle to the newly constructed resource

25               On failure this may be set to a ‘null’ handle

26       **UnBuild(composite-resource-handle) -> resource-handle[]**

27           The interface is used to disassemble a composite resource into the set of base

resources used to build it. It is used when resources are to be returned to the BRDS.

Inputs:

Composite-resource-handle, handle of the composite resource.

Outputs:

Resource-handle[], array of handles to the base resources

## **Managed Resource Service**

As described earlier, an MRS for a resource type (composite or base) encapsulates the knowledge of the operation of the resource. An MRS is instantiated for a resource when the resource is created by the corresponding PRS. A reference to the MRS instance is returned from the build operation on the PRS.

Recall that an MRS for a composite resource type can use the MRSs of the sub-resources in the management of the composite resource. This structure is shown in Figure 14 for the web site computing environment of Figure 4. Every computing environment that has resources allocated to it has one designated **root MRS**. The root MRS in Figure 14 is the Web site MRS 401. All other MRS instances are associated beneath this root MRS. Only the root MRS communicates with the BRDS 219. The root MRS generates resource management requests 1405 based on the system's knowledge taking into account factors such as the system's performance and health. These requests may modify the set of base resources used by computing environment corresponding to the root MRS. An MRS must implement the operations described in the BRDS resource management interface described in cross-referenced application "Hierarchical Resource Management for a Computing Utility", having docket number YOR920030587PCT1. Lower level MRSs can only generate internal requests for more capacity that will be propagated through the hierarchy to the root MRS. Note that an MRS for a given resource type may be a root MRS in one computing environment but not in another.

To compose MRSs, the resource management interface is replicated at each level, with upper level interfaces receiving and propagating requests between BRDS and the lowest level MRSs, such as between web server MRS 493 and web site MRS 401. Intermediate MRSs may filter out requests when appropriate. For example, web site MRS 401 may choose to respond to a resource request from a lower level MRS by reconfiguring existing resources instead of requesting resources from the BRDS. The reasoning behind this structure is that a high level MRS can make more intelligent decisions based on knowledge of the entire composite including interdependencies between the requesting sub resource and other resources.

Third party workload management systems, such as eWLM which is described in an application entitled "SYSTEM AND METHOD FOR ADAPTIVE ADMISSION CONTROL AND RESOURCE MANAGEMENT FOR SERVICE TIME GUARANTEES", having docket number YOR920020201US1, can also be represented in the PMRS structure as shown in Figure 15. Two computing environments 1501 and 1503 employ resources such as servers 1505, and workload management software 1507 to manage workload over the servers. The workload management software is configured and build into the composite resource for the computing environment by the PRS part of PMRS 1509. The MRS part of PMRS 1509 functions as an adapter capable of interacting with the workload management software. This adapter layer interprets its state and requests more resources if necessary from BRDS 219.

### **Internal Structure of MRS**

The functionality and operational interface provided by an MRS is resource specific. A wide variety of functions are possible. In general, all of the functions contained within an MRS are associated with the management of a resource which includes configuration, operation, and maintenance. For example, an MRS may be capable of monitoring, which includes gathering events or other metrics from other sources in the system, applying some function such as smoothing or aggregation, and generating metrics events that describe the operational state of the

1 resource. A generic interface may be provided to indicate which metrics or events are produced  
2 and consumed by an MRS. The metrics collected within an MRS may be reported to another part  
3 of the system such as a billing system or another MRS. Another type of functionality is  
4 configuration, and generic operations may be provided to get and set configuration parameters,  
5 and to change the state of a resource to optimize its behavior. A third type of functionality  
6 provides self-managing capabilities by receiving metrics or events that describe the state of the  
7 managed resource and determining what action to take, if any, to keep the resource operating at  
8 the required level (e.g., as specified in an SLA). There are essentially two types of operations  
9 that can be initiated in this way: changing the operational state of the resource by reconfiguring it  
10 or its sub-resources using the appropriate configuration operations, and changing the capacity of  
11 a resource or sub-resource. Changing the configuration of a sub-resource can be done through its  
12 MRS. Changing capacity of resources is a more involved process. An increase in capacity could  
13 be requested because of failure to meet an SLA or because of some projected future load  
14 modification. Either the requested new resources would have to be specified on the request or an  
15 analysis would have to be done to determine where the bottleneck is and what resources (or  
16 configurations) would be required to alleviate the bottleneck. Once this is done a realizer service  
17 is used to obtain a shopping list of basic resources, and requests are sent to the BRDS to obtain  
18 the resources. Once the resources are obtained, the builder is invoked with the resource handles.  
19 A decrease in capacity could also be caused by an SLA and would not require a realization step  
20 because existing resources would be returned.

## 21 **Process and Methodology to Incorporate New Resources or Services**

22 One of the benefits of the invention is the ability of the apparatus to cope with a dynamic system  
23 in which new resource types and services are introduced. Figure 16 shows a process to be  
24 followed when a new base resource is introduced. The process begins in step 1601 with a  
25 determination of whether or not the resource to be introduced is a base resource. If the resource  
26 is a base resource, flow passes to step 1613, where a BRLS for the new base resource type needs  
27 to be provided. The BRLS is then activated in step 1615, and updated in step 1617 with the



information on available resource instances through the BRLSs administrative interfaces (see cross-referenced application YOR920030587PCT1). Next in step 1619 a PRS implementation must be provided for the new resource. The new PRS is activated in step 1621, as there must be a single PRS for every resource type in the system. In step 1623, an MRS implementation must be provided for the new resource. Finally in step 1611, the logic of any component that wishes to leverage the new resource needs to be updated. It is this step that allows preexisting services to take advantage of the newly defined resource. The circumstances which a new resource is incorporated into preexisting services, and the extent to which this incorporation is automated, can be a matter of service provider or installation policy. Examples of components that need to be updated are PRSs and MRSs of composite resources that can now realize the resource they represent using the new base resource. For example, introducing a new type of machine (such as an xSeries server), a Web server PRS may need to be updated to offer realizations which includes also the possibility of using the new machine type. In particular, both the realizer function of a PRS and the build function may need to be updated. No changes are necessary in other parts of the system.

For a composite resource, the process continues with step 1603, modeling the composite resource's possible structures; list of sub resources, how are they interrelated and in particular how they are configured to work together. The result is step 1605, an implementation of a realizer and builder functions in a PRS for the new composite resource that uses existing resources' PRSs as part of its logic to realize and build the composite resource. In step 1607 the new PRS is activated. Next in step 1609, an MRS implementation is provided for the new composite resource that uses existing sub resources' MRSs as part of its logic to manage the composite resource. The process completes with updating existing higher level composite resources where necessary to take advantage of the new composite resource in step 1611. For example, a Web site composite resource may be updated to take advantage of a new Secure Web server composite resource (assuming that previously only simple Web server composite resource was available).

1 The process to incorporate and support a new service type is similar to the one for composite  
2 resources as services are modeled as composite resources. These processes enable a service  
3 provider or other agent involved with the operation of the computing utility to increase the  
4 services offered to customers. Note that the implementer of the new services may be distinct  
5 from the service provider or operator of the computing utility; new implementations of resources  
6 may be provided as packaged upgrades to the service provider, who then installs them to add new  
7 resources. A service provider in this context could represent entities either external or internal to  
8 a customer or customer's organization. An example of the latter is an IT department within a  
9 company. Other steps that may be necessary for a service beyond what is needed for an  
10 internally used composite resource, such as publishing the service to the public, are beyond the  
11 scope of this invention.

## 12 Use Cases

13 Two use cases are shown in the Figures to illustrate the componentized operation of provisioned  
14 and managed resources. The first use case, in Figure 17, shows a computing environment  
15 consisting of a set of servers. The set is modeled as a composite resource for which a PRS 1701  
16 and an MRS 1703 exist. The MRS 1703 is the designated root MRS for the computing  
17 environment. Each individual server is modeled as a resource, with PRS 1705. There are three  
18 servers in the set, as exemplified by MRS 1707. Server set MRS receives monitoring data 1709  
19 that indicates that the capacity of the group should be increased. MRS 1703 decides to add a  
20 server. It issues realize request 1713 to server PRS 1705 to obtain a resource list B and plan to  
21 build another server. Only B is shown in the figure. It then issues a resource request 1715 to the  
22 BRDS 219 to obtain the resources. In this case there are no advance reservations; the BRDS  
23 returns resource handles. MRS 1703 then invokes the builder of server PRS 1705 via request  
24 1719 with the resource handles. A new server MRS 1721 is instantiated. Upon completion of the  
25 build request, MRS 1703 may update its state with the new server's resource handle and  
26 additionally perform any configuration needed to add the new server to the set.

1 The second use case, in Figure 18, shows a more complex resource structure in the form of a  
2 capacity increase for a Web site. The root MRS is the web site MRS 401. Again, monitoring  
3 data 1803 arrives at the root MRS that indicates that the capacity of the Web site must be  
4 increased. MRS 401 decides to add a Web server. It issues a realize request 1807 to the Web  
5 server PRS 503. The Web server MRS in turn invokes the realize operation 1811 of the server  
6 PRS 507. The list of base resources is obtained and a request 1815 is issued to BRDS 219 as in  
7 the previous use case. The resource handles are returned and passed to the builders of the  
8 component resources via build requests 1819 and 1821. The server and Web server builders  
9 instantiate MRSs 1825 and 1827 to represent the new server and Web server instances,  
10 respectively. Finally, the root MRS 401 configures the new server via configure request 1831 to  
11 the load balancer through its MRS 405 to include the new server as part of the Web site.

12 Thus, the present invention includes a method for provisioning one or more computing  
13 environment in a computing utility, said method comprising: employing a description of each of  
14 the one or more computing environment, the description directing an invocation of a first service  
15 which performs the provisioning, the description being a composite resource definition; and  
16 invoking the first service to perform the provisioning.

17 In some embodiments of the method, the step of obtaining a plan, or realization, is repeated when  
18 there are changes in the description of the computing environment, changes in the set of base or  
19 composite resource types, or any combination of these changes.

20 In some embodiments of the method, the steps of selecting a plan and requesting the resources  
21 associated with it are repeated when there are changes in the selection criteria, resource  
22 availability, resource reservation status, changes in plans, or any combination of these changes.

23 In further embodiments of the present invention, the mapping of provisioning resource services  
24 and managed resource services to resource types and instances may vary based on considerations  
25 such as performance, availability, and security. For example, multiple provisioning resource

1 services may be used to supply instances of a resource type for better performance. A single  
2 provisioning resource service that accepts the type as input may be used for provisioning all  
3 resource types. A single managed resource service may manage all instances of a resource type.  
4 Alternatively, if isolation is required between components of different computing environments,  
5 a single managed resource service may manage multiple instances of a resource type within a  
6 computing environment. Further, the partitioning of provisioning and managed resource services  
7 into two components can also vary. For example, they can be realized as a single service or a  
8 collection of services. Finally, provisioning resource services and managed resource services  
9 may themselves be realized as single services or collections of services. For example, the builder  
10 and realizer functions of the provisioning resource service could be instantiated as separate  
11 services. Those skilled in the art will realize that many such variations are possible, and that  
12 these embodiments are merely illustrative and not restrictive of the invention.

13 The invention also includes an apparatus comprising a set of provisioning resource services that  
14 transform a description of a computing environment into a plan for constructing a set of  
15 resources to provide the computing environment, where the computing environment is  
16 represented as a composite resource definition. In some embodiments of the apparatus, the set  
17 of provisioning resource services is to construct the resources into the computing environment  
18 according to the plan, and the one or more provisioning resource service to make available a  
19 managed resource service to manage the computing environment.

20 The invention also includes a method that obtains a description of a computing environment,  
21 wherein the description is represented as a composite resource definition, and transforms a  
22 description of a computing environment into a plan for constructing a set of resources to provide  
23 the computing environment. In some embodiments of the method, it also constructs the  
24 resources into the computing environment according to the plan, and makes available a managed  
25 resource service to manage the computing environment.

26 Variations described for the present invention can be realized in any combination desirable for

1 each particular application. Thus particular limitations, and/or embodiment enhancements  
2 described herein, which may have particular advantages to a particular application need not be  
3 used for all applications. Also, not all limitations need be implemented in methods, systems  
4 and/or apparatus including one or more concepts of the present invention.

5 The present invention can be realized in hardware, software, or a combination of hardware and  
6 software. A visualization tool according to the present invention can be realized in a centralized  
7 fashion in one computer system, or in a distributed fashion where different elements are spread  
8 across several interconnected computer systems. Any kind of computer system - or other  
9 apparatus adapted for carrying out the methods and/or functions described herein - is suitable. A  
10 typical combination of hardware and software could be a general purpose computer system with  
11 a computer program that, when being loaded and executed, controls the computer system such  
12 that it carries out the methods described herein. The present invention can also be embedded in a  
13 computer program product, which comprises all the features enabling the implementation of the  
14 methods described herein, and which - when loaded in a computer system - is able to carry out  
15 these methods.

16 Computer program means or computer program in the present context include any expression, in  
17 any language, code or notation, of a set of instructions intended to cause a system having an  
18 information processing capability to perform a particular function either directly or after  
19 conversion to another language, code or notation, and/or reproduction in a different material  
20 form.

21 Thus the invention includes an article of manufacture which comprises a computer usable  
22 medium having computer readable program code means embodied therein for causing a function  
23 described above. The computer readable program code means in the article of manufacture  
24 comprises computer readable program code means for causing a computer to effect the steps of a  
25 method of this invention. Similarly, the present invention may be implemented as a computer  
26 program product comprising a computer usable medium having computer readable program code

1 means embodied therein for causing a a function described above. The computer readable  
2 program code means in the computer program product comprising computer readable program  
3 code means for causing a computer to effect one or more functions of this invention.

4 Furthermore, the present invention may be implemented as a program storage device readable by  
5 machine, tangibly embodying a program of instructions executable by the machine to perform  
6 method steps for causing one or more functions of this invention.

7 It is noted that the foregoing has outlined some of the more pertinent objects and embodiments of  
8 the present invention. This invention may be used for many applications. Thus, although the  
9 description is made for particular arrangements and methods, the intent and concept of the  
10 invention is suitable and applicable to other arrangements and applications. It will be clear to  
11 those skilled in the art that modifications to the disclosed embodiments can be effected without  
12 departing from the spirit and scope of the invention. The described embodiments ought to be  
13 construed to be merely illustrative of some of the more prominent features and applications of the  
14 invention. Other beneficial results can be realized by applying the disclosed invention in a  
15 different manner or modifying the invention in ways known to those familiar with the art.

1 CLAIMS

2 We claim:

3 1. A method for provisioning at least one computing environment in a computing utility, said  
4 method comprising:

5 employing a description of each of said at least one computing environment, said description  
6 directing an invocation of a first service which performs said provisioning, said description being  
7 a composite resource definition; and

8 invoking said first service to perform said provisioning.

9 2. A method as recited in claim 1, wherein said first service is a provisioning resource service,  
10 and said method further comprises said first service making available a managed resource service  
11 for managing each of said at least one computing environment.

12 3. A method as recited in claim 2, wherein there is at least one provisioning resource service  
13 associated with each type of base resource and each type of composite resource.

14 4. A method as recited in claim 2, wherein said provisioning resource service provides a plurality  
15 of provisioning instances.

16 5. A method as recited in claim 2, wherein the step of invoking comprises:  
17 obtaining at least one plan to satisfy said at least one computing environment;  
18 selecting a particular plan;

requesting base resources corresponding to said particular plan from said computing utility;  
repeating the steps of selecting and requesting until said base resources corresponding to said particular plan is obtained; and  
configuring said base resources into an instance of said computing environment.

6. A method as recited in claim 5, wherein said base resources satisfy a request criterion taken from a group of request criteria comprising:

a designated time at which the resources will be provisioned;

duration of resource availability;

resource-specific selection criteria; and

any combination of these request criteria.

7. A method as recited in claim 2, wherein said provisioning resource service is comprised of a group of provisioning resource services, each member of said group of provisioning resource services is associated with a particular type of resource in said composite resource description.

8. A method as recited in claim 7, wherein the step of invoking further comprising employing said provisioning resource service for the computing environment by recursively employing each member of said group according to said composite resource description.

9. A method as recited in claim 2, wherein each base resource has a base resource instance, and wherein there is at least one managed resource service associated with each base resource instance.



1 10. A method as recited in claim 2, wherein each composite resource has a composite resource  
2 instance, and wherein there is at least one managed resource service associated with each  
3 composite resource instance.

4 11. A method as recited in claim 2, wherein said managed resource service is comprised of a  
5 group of managed resource services, each member of said group of managed resource services is  
6 associated with a particular instance of resource in said composite resource description.

7 12. A method as recited in claim 5, wherein said provisioning resource service is comprised of a  
8 group of provisioning resource services, each member of said group of provisioning resource  
9 services is associated with a particular type of resource in said composite resource description;  
10 and

11 wherein the step of invoking further comprising employing said provisioning resource service for  
12 the computing environment by recursively employing each member of said group according to  
13 said composite resource description; and

14 wherein the step of obtaining said at least one plan comprises recursively obtaining at least one  
15 plan for each sub-resource, and combining the result; and

16 wherein the step of configuring said base resources into an instance of said computing  
17 environment comprises recursively configuring base resources corresponding to each  
18 sub-resource, and configuring the result.

19 13. A method as recited in claim 5, wherein the step of obtaining further comprises specifying at  
20 least one condition restricting the plans obtained.

21 14. A method as recited in claim 5, wherein the step of selecting said plan satisfies a selection

1 criterion taken from a group of selection criteria comprising:

2 cost;

3 performance; and

4 number of preferences satisfied; and

5 any combination of these selection criteria.

6 15. A method as recited in claim 5, further comprising transforming reservations for base  
7 resources into allocated base resources.

8 16. A method as recited in claim 5, wherein the step of obtaining is repeated according to a  
9 repetition criterion taken from a group of repetition criteria comprising:

10 changes in the description of the computing environment;

11 changes in the set of base resource types;

12 changes in the set of composite resource types; and

13 any combination of these changes.

14 17. A method as recited in claim 5, wherein the step of selecting and requesting are repeated  
15 according to a repetition criterion taken from a group of repetition criteria comprising:

16 changes in the selection criteria;

changes in resource availability;

changes in resource reservation status;

changes in said at least one plan; and

any combination of these changes.

18. A method as recited in claim 15, further comprising iterating the steps of repeating and transforming until said base resources are obtained.

19. A method as recited in claim 5, wherein the steps of obtaining, selecting, requesting, and configuring are invoked at least once by at least one managed resource service to add resources to said computing environment.

20. An article of manufacture comprising a computer usable medium having computer readable program code means embodied therein for causing provisioning of at least one computing environment in a computing utility, the computer readable program code means in said article of manufacture comprising computer readable program code means for causing a computer to effect the steps of claim 1.

21. A program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform method steps for provisioning at least one computing environment in a computing utility, said method steps comprising the steps of claim 1.

22. A method comprising adding a base resource type to an operating computing utility, said step of adding comprising:

1       furnishing a base resource library service to represent any instance of said base resource  
2       type;

3       adding any instance of said base resource type to the base resource library service;

4       developing an implementation of a provisioning resource service for said base resource  
5       type;

6       activating said provisioning resource service; and

7       creating an implementation of a managed resource service for said base resource type.

8       23. A method as recited in claim 22, further comprising:

9       updating any provisioning resource service of any composite resource type to use said  
10      provisioning resource service for said base resource type; and

11      updating any managed resource service of any composite resource type to use said  
12      managed resource service for said base resource type.

13      24. A method as recited in claim 22, wherein the step of adding a base resource type is  
14      performed by a service provider to increase service offerings to at least one customer of said  
15      service provider.

16      25. A method comprising adding a composite resource type to an operating computing utility,  
17      said step of adding comprising:

18      creating an implementation of a provisioning resource service for said composite resource type,  
19      said provisioning resource service using the provisioning resource services for sub-resources of

1     said composite resource type;

2     activating said provisioning resource service; and

3     developing an implementation of a managed resource service for said composite resource type,

4     said managed resource service using the managed resource services for sub-resources of said

5     composite resource type.

6     26. A method as recited in claim 25, further comprising:

7     updating any provisioning resource service for a composite resource to use said provisioning

8     resource service for said composite resource type; and

9     updating any managed resource service for a composite resource to use said managed resource

10    service for said composite resource type.

11    27. A method as recited in claim 25, wherein the step of adding a composite resource type is

12    performed by a service provider to increase service offerings to at least one customer of said

13    service provider.

14    28. A method comprising adapting at least one third party workload management system for a

15    computing environment to a computing utility, said step of adapting comprising:

16    developing a managed resource service, said managed resource service interpreting a state of said

17    at least one third party workload management system; and

18    modifying a set of base resources comprising said computing environment as indicated by the

19    state of said at least one third party workload management system.

1     29. A method as recited in claim 2, further comprising:

2     receiving an event for returning at least one base resource from said at least one base resource  
3     from at least one computing environment of said at least one computing environment;

4     unconfiguring said at least one base resource from at least one computing environment of said at  
5     least one computing environment;

6     returning said at least one base resource to said computing utility; and

7     making available said at least one base resource for allocation to another computing  
8     environment.

9     30. A method as recited in claim 9, wherein there is one managed resource service for all  
10    instances of a resource type within a single computing environment.

11    31. Method as recited in claim 1, wherein said description further comprises at least one  
12    condition on said composite resource definition.

13    32. A method as recited in claim 1, further comprising receiving a customer request for at least  
14    one of said at least one computing environment.

15    33. A method as recited in claim 1, wherein said composite resource definition is a recursive  
16    combination of resources, said combination being a combination taken from a group of  
17    combinations consisting of:

18         at least one base resource;

19         at least one other composite resource; and

20         any combination of base resources and composite resources.

34. A method as recited in claim 11, wherein said managed resource service employs at least one operation available to a particular member of said group of managed resource services, said operation being performed by at least one member of said group of managed resource services, said at least one operation being taken from a group of operations comprising:

delivering at least one metric based on monitored data from said computing environment;

delivering at least one event based on monitored data from said computing environment;

modifying a state of said computing environment;

modifying at least one resource of said computing environment;

invoking at least one particular operation to modify the computing environment in response to at least one particular metric; and

invoking at least one particular operation to modify the computing environment in response to at least one particular event.

35. An article of manufacture comprising a computer usable medium having computer readable program code means embodied therein for causing addition of a base resource type to an operating computing utility, the computer readable program code means in said article of manufacture comprising computer readable program code means for causing a computer to effect the steps of claim 1.

36. An apparatus to provision at least one computing environment in a computing utility, said apparatus comprising:

a first processor module to employ a description of each of said at least one computing

environment, said description directing an invocation of a first service which performs said provisioning, said description being a composite resource definition; and

a second processor module to invoke said first service to perform said provisioning.

37. An apparatus as recited in claim 36, wherein said second processor module makes available a managed resource service to manage each of said at least one computing environment.

38. A computer program product comprising a computer usable medium having computer readable program code means embodied therein for causing provisioning of at least one computing environment in a computing utility, the computer readable program code means in said computer program product comprising computer readable program code means for causing a computer to effect the functions of claim 36.

39. A program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform method steps for adding a base resource type to an operating computing utility, said method steps comprising the steps of claim 1.

40. An article of manufacture comprising a computer usable medium having computer readable program code means embodied therein for causing addition of a composite resource type to an operating computing utility, the computer readable program code means in said article of manufacture comprising computer readable program code means for causing a computer to effect the steps of claim 1.

41. A program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform method steps for adding a composite resource type to an operating computing utility, said method steps comprising the steps of claim 1.

42. An apparatus comprising:



1 a plurality of provisioning resource services to transform a description of a computing  
2 environment into a plan for constructing a set of resources to provide said computing  
3 environment, wherein the computing environment is represented as a composite resource  
4 definition.

5 43. An apparatus as recited in claim 42, wherein said plurality of provisioning resource services  
6 is to construct said resources into said computing environment according to said plan, and said  
7 at least one provisioning resource service to make available a managed resource service to  
8 manage the computing environment.

9 44. A method comprising:

10 obtaining a description of a computing environment, wherein said description is represented as a  
11 composite resource definition; and

12 transforming a description of a computing environment into a plan for constructing a set of  
13 resources to provide said computing environment.

14 45. An method as recited in claim 44, further comprising:

15 constructing said resources into said computing environment according to said plan; and

16 making available a managed resource service to manage the computing environment.

17 46. An article of manufacture comprising a computer usable medium having computer readable  
18 program code means embodied therein for obtaining a description of a computing environment,  
19 the computer readable program code means in said article of manufacture comprising computer  
20 readable program code means for causing a computer to effect the steps of claim 44.

1 47. A program storage device readable by machine, tangibly embodying a program of  
2 instructions executable by the machine to perform method steps for obtaining a description of a  
3 computing environment, said method steps comprising the steps of claim 44.

4 48. An article of manufacture comprising a computer usable medium having computer readable  
5 program code means embodied therein for causing adaptation of at least one third party workload  
6 management system for a computing environment to a computing utility, the computer readable  
7 program code means in said article of manufacture comprising computer readable program code  
8 means for causing a computer to effect the steps of claim 28.

9 49. A program storage device readable by machine, tangibly embodying a program of  
10 instructions executable by the machine to perform method steps for adapting at least one third  
11 party workload management system for a computing environment to a computing utility, said  
12 method steps comprising the steps of claim 28.

13 50. A computer program product comprising a computer usable medium having computer  
14 readable program code means embodied therein for causing development of a plan, the computer  
15 readable program code means in said computer program product comprising computer readable  
16 program code means for causing a computer to effect the functions of claim 42.

17 51. An method as recited in claim 1, further comprising a Provisioned and Managed Resource  
18 Service of a composite resource using the Provisioned and Managed Resource Service of a  
19 sub-resource to create and manage instances of the sub-resource.

1           **COMPONENTIZED AUTOMATIC PROVISIONING AND MANAGEMENT OF**  
2           **COMPUTING ENVIRONMENTS FOR COMPUTING UTILITIES**

3                           **ABSTRACT**

4       The present invention provides systems, methods and apparatus for automatically provisioning  
5       and managing resources in a computing utility. Its automation procedures are based on a  
6       resource model which allows resource-specific provisioning and management tasks to be  
7       encapsulated into components for reuse. These components are assembled into more complex  
8       structures and finally computing services. This invention provides a method for constructing a  
9       computing service from a set of resources given a high level specification. Once constructed, the  
10      service includes a component that provides management function, which can allow modification  
11      of its underlying set of resources.

Fig. 1




Fig. 2




Fig. 2a




Fig. 3




Fig. 4






Fig. 5




Fig. 6





Fig. 7



Fi. 8




Fig. 9




Fig. 10

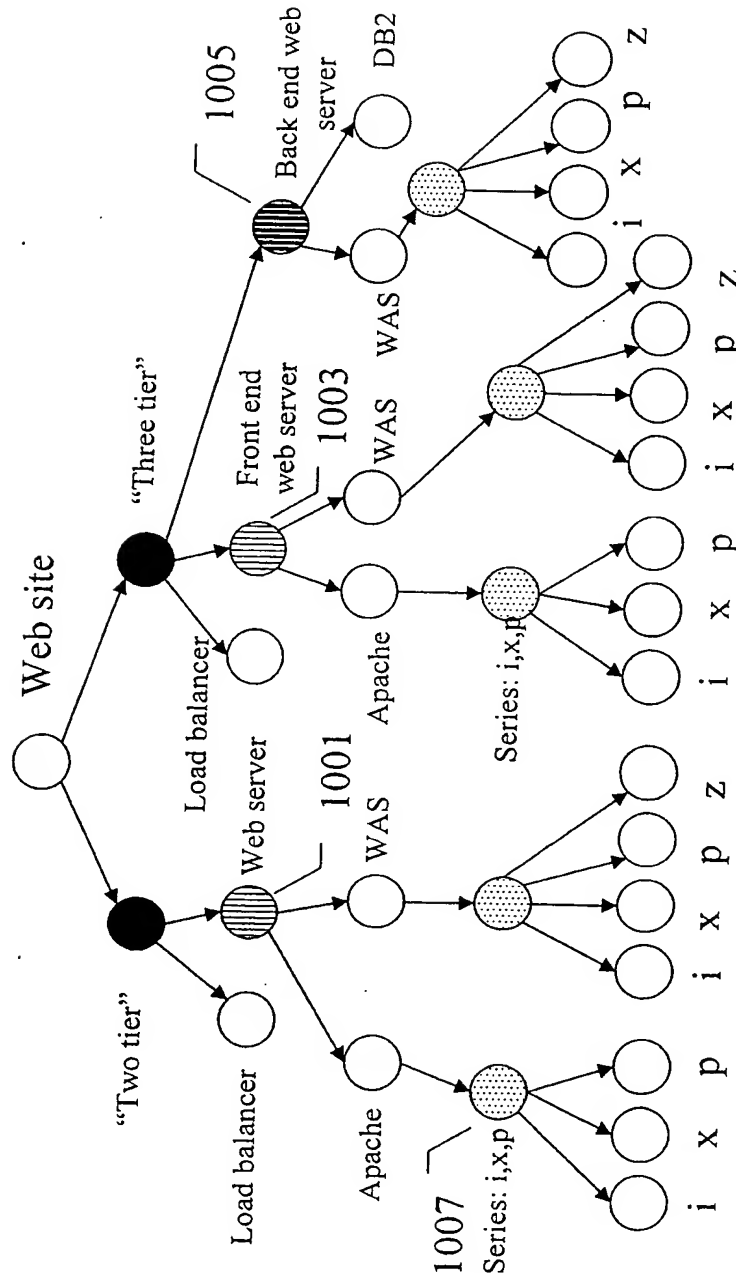


Fig. 11




Fig. 12

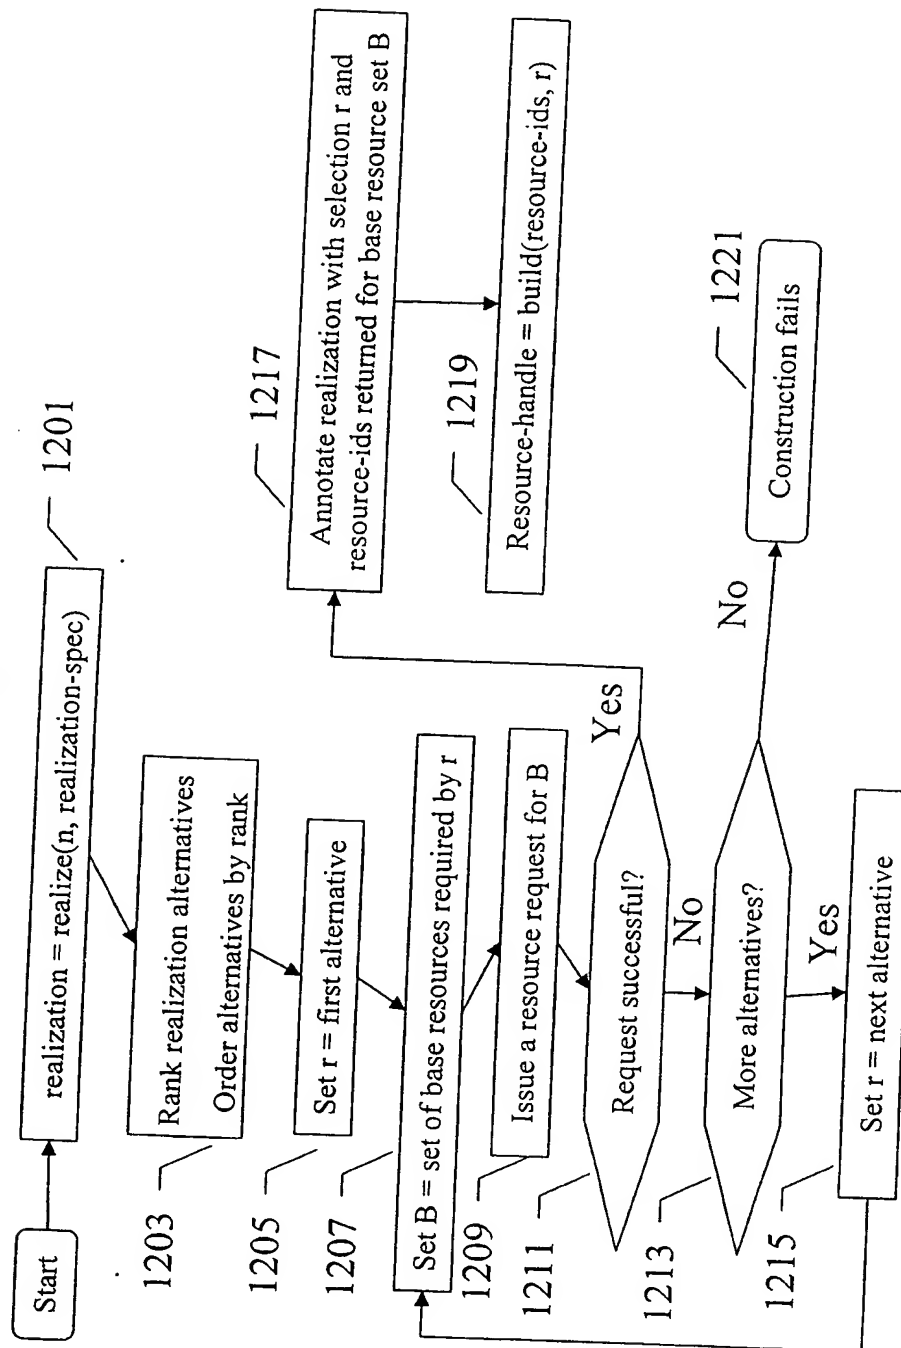




Fig. 13

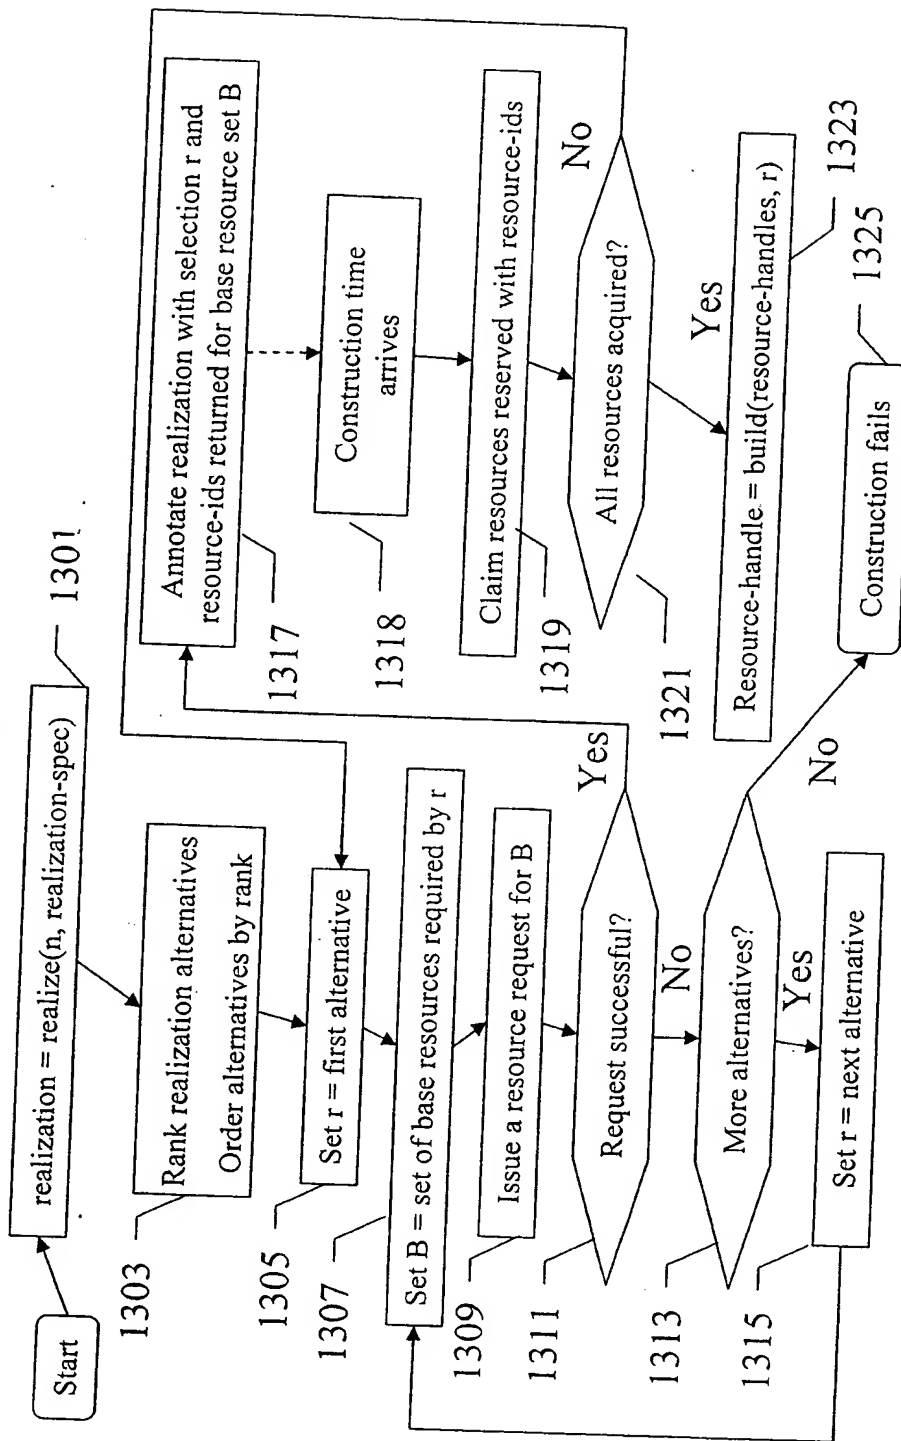


Fig. 14




Fig. 15




Fig. 16

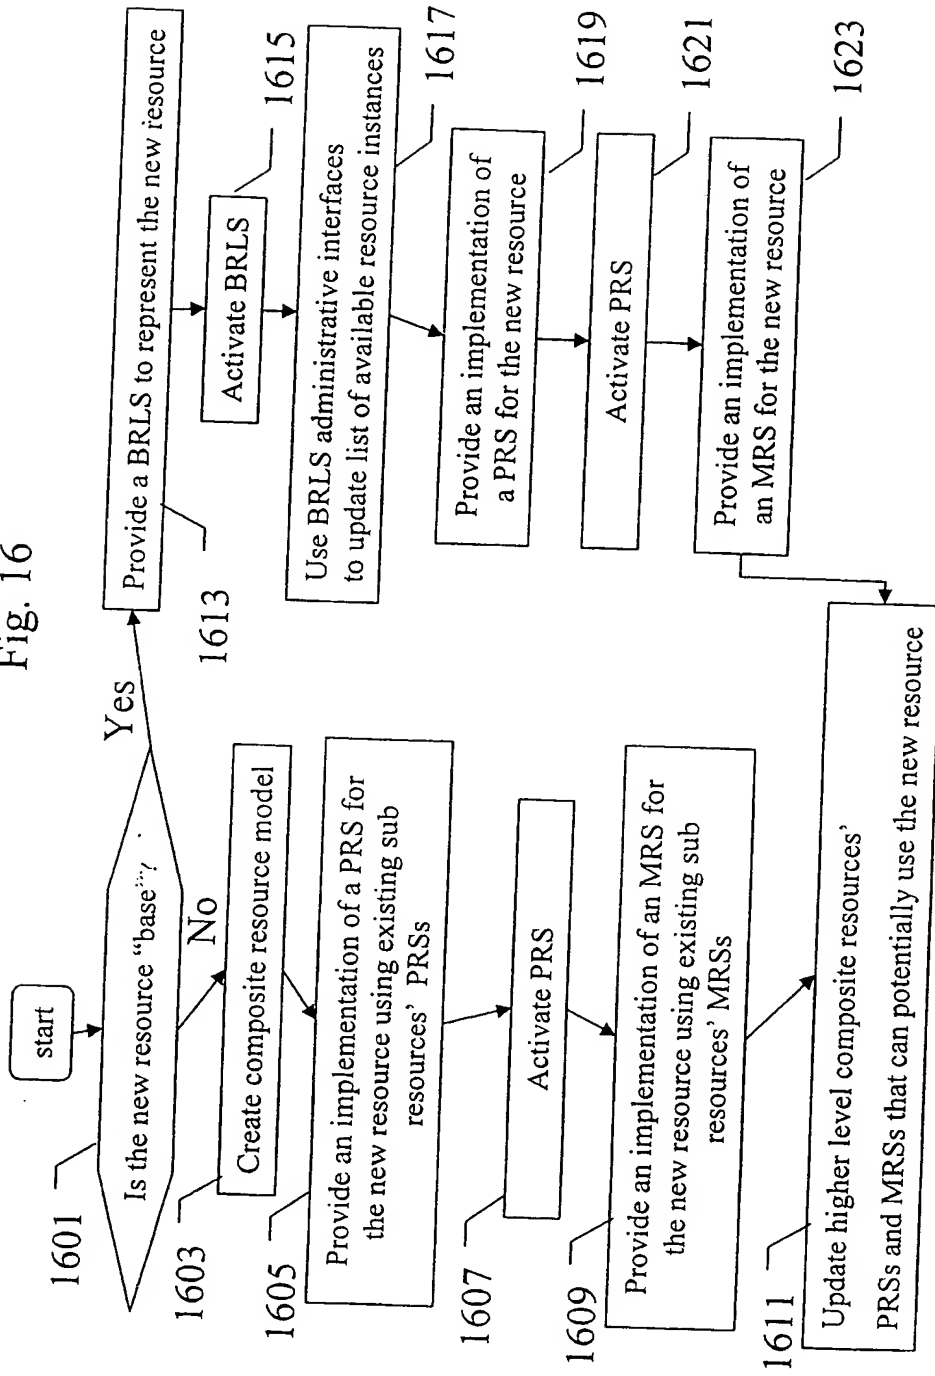


Fig. 17

